

What is claimed is:

- [Claim 1]** 1. A method for versioning in a storage architecture that manages node ranges, said method comprising:
- a. receiving a node modification request from a database system;
 - b. copying, to a storage, a node range to which said modification is to be made;
 - c. labeling said copied node range with an identifier; and
- wherein said labeled node range is locatable via said identifier and a hash on said node range.
- [Claim 2]** 2. A method as per claim 1, wherein said identifier is any of the following: a timestamp or a LSN.
- [Claim 3]** 3. A method as per claim 1, wherein said storage is a transient storage.
- [Claim 4]** 4. A method as per claim 1, wherein said node modification request is any of the following: a node insertion request, a node update request, or a node deletion request.
- [Claim 5]** 5. A method as per claim 1, wherein said method is implemented across networks.
- [Claim 6]** 6. A method as per claim 5, wherein said network is any of the following: local area network, wide area network, or the Internet.

[Claim 7] 7. A method as per claim 1, wherein said node ranges are associated with hierarchical node data that is derived from any of: a structured document, a computer network, or a directory file system.

[Claim 8] 8. A method as per claim 7, wherein said structured document is an XML document.

[Claim 9] 9. A method for versioning in a storage architecture that manages node ranges via a node id range index, said each node assigned a node id value and a set of nodes forming a node range, each entry in said node id range index pointing to a node range and its range identifier, RID, said method comprising:

- a. receiving a node modification request for a range;
- b. shadowing nodes in said range to a Version Hash Table based on RID;
- c. assigning a time identifier to copies of said range;

wherein a node in said shadowed range is locatable via said time identifier and RIDs.

[Claim 10] 10. A method as per claim 9, wherein said time identifier is any of the following: timestamp or LSN.

[Claim 11] 11. A method as per claim 9, wherein new readers, after a modification, access current nodes through a new RID.

[Claim 12] 12. A method as per claim 9, wherein previous readers access old nodes via the same RID and hashing the same RID to locate the shadowed copy in said Version Hash Table.

[Claim 13] 13. A method as per claim 9, wherein when modifications cause nodes in a range to be moved to a new RID, previous readers are redirected from the new RID to the old RID via a Redirection Hash Table.

[Claim 14] 14. A method as per claim 9, wherein when modifications cause nodes in a range to be moved to a new RID, previous readers are redirected from the new RID to the old RID via an index that describes where old versions are in said Version Hash Table.

[Claim 15] 15. A method as per claim 9, wherein said shadowed nodes are copied to a transient storage.

[Claim 16] 16. A method as per claim 9, wherein said method is implemented across networks.

[Claim 17] 17. A method as per claim 16, wherein said network is any of the following: local area network, wide area network, or the Internet.

[Claim 18] 18. A method as per claim 9, wherein, for range deletions, the range being deleted is moved to reserved RID RIDFF.

[Claim 19] 19. A method as per claim 18, wherein a reader hashes said Redirection Hash Table on RIDFF to find the correct Version Hash Table entry.

[Claim 20] 20. A method as per claim 9, wherein said node ranges are associated with hierarchical node data that is derived from any of: a structured document, a computer network, or a directory file system.

[Claim 21] 21. A method as per claim 20, wherein said structured document is an XML document.

[Claim 22] 22. A method as per claim 9, wherein said node modification request is any of the following: a node insertion request, a node update request, or a node deletion request.

[Claim 23] 23. A method as per claim 9, wherein said method favors new readers over old readers as old readers have to redirected to shadowed storage.

[Claim 24] 24. An article of manufacture comprising computer readable program code implementing a method for transient versioning in a storage architecture that manages node ranges via a node id range index, said each node assigned a node id value and a set of nodes forming a node range, each entry in said node id range index pointing to a node range and its range identifier, RID, said method comprising:

- a. computer readable program code identifying a node modification request for a range;
 - b. computer readable program code shadowing nodes in said range to a Version Hash Table based on RID;
 - c. computer readable program code assigning a time identifier to copies of said range;
- wherein a node in said shadowed range is locatable via said time identifier and RIDs.

[Claim 25] 25. An article of manufacture comprising computer readable program code implementing a method for versioning in a storage architecture that manages node ranges, said method comprising:

- a. computer readable program code identifying a request for node modification from a database system;
 - b. computer readable program code copying, to a storage, a node range to which said modification is to be made;
 - c. computer readable program code labeling said copied node range with an identifier; and
- wherein said labeled node range is locatable via said identifier and a hash on said node range.